arXiv:cs/0509071v1 [cs.GT] 22 Sep 2005

# CP-nets and Nash equilibria

Krzysztof R. Apt[1,2,3], Francesca Rossi[4], and Kristen Brent Venable[4]
[1] School of Computing, National University of Singapore
[2] CWI, Amsterdam
[3]University of Amsterdam, the Netherlands
[4]Department of Pure and Applied Mathematics, University of Padova

**Abstract**

*We relate here two formalisms that are used for different purposes in reasoning about multi-agent systems. One of them are strategic games that are used to capture the idea that agents interact with each other while pursuing their own interest. The other are CP-nets that were introduced to express qualitative and conditional preferences of the users and which aim at facilitating the process of preference elicitation.*

*To relate these two formalisms we introduce a natural, qualitative, extension of the notion of a strategic game. We show then that the optimal outcomes of a CP-net are exactly the Nash equilibria of an appropriately defined strategic game in the above sense. This allows us to use the techniques of game theory to search for optimal outcomes of CP-nets and vice-versa, to use techniques developed for CP-nets to search for Nash equilibria of the considered games.*

## 1 Introduction

One of the main tools in the area of multi-agent systems is game theory, notably strategic games. They formalize in a simple and powerful way the idea that agents interact with each other while pursuing their own interest. The interaction is captured by the fact that actions (strategies) are taken simultaneously, while agents' interests are expressed by means of the utility (payoff) function that each agent wishes to maximize.

*CP-nets* (Conditional Preference nets) are an elegant formalism for representing conditional and qualitative preferences, see [3, 2]. They model such preferences under a *ceteris paribus* (that is, 'all else being equal') assumption. The CP-net represents a complex 'joint preference distribution' in a compact form. Preference elicitation in such a framework appears to be natural and intuitive.

Research on CP-nets focused on its modeling capabilities and algorithms for solving various natural problems related to their use. Also, computational complexity of these problems was extensively studied. An outcome of a CP-net is an assignment of values to its variables. One of the fundamental problems is that of finding an optimal outcome, i.e., the one that cannot be improved in presence of the adopted preference statements. This is in general a complex problem since it was found that finding optimal outcomes and testing for their existence is NP-hard in general. In contrast, for so-called acyclic CP-nets this is an easy problem which can be solved by a linear time algorithm.

The aim of this paper to show the relationship between CP-nets and game theory, and explain how game-theoretic techniques developed for the analysis of strategic games can be fruitfully used to study CP-nets. To this end, we introduce a generalization of the customary strategic games (see, e.g., [5],) in which each player has to his disposal a strict preference relation on his set of strategies, parametrized by a joint strategy of his opponents. We call such games *strategic games with parametrized preferences*.

The cornerstone of our approach are two results closely relating CP-nets to such games. They show that the optimal outcomes of a CP-net are exactly the Nash equilibria of an appropriately defined strategic game with parametrized preferences. This allows us to transfer techniques of game theory to CP-nets.

To find Nash equilibria in strategic games, reduction techniques have been studied which reduce the game by eliminating some players' strategies, thus obtaining a smaller game. We introduce two counterparts of such game-theoretic techniques that allow us to reduce a CP-net while maintaining its optimal outcomes. We also introduce a method of simplifying a CP-net by eliminating so-called redundant variables from the variables parent sets. Both techniques simplify the search for optimal outcomes of a CP-net.

The paper is organized as follows. Section 2 pro-

vides the basic definitions of CP-nets. Then, Section 3 introduces our generalized notion of games, Section 4 shows how to pass from a CP-net to a game, and Section 5 handles the opposite direction. Then, Section 6 introduces the concept of reduced CP-nets, and Section 7 shows how to exploit techniques developed in games for CP-nets. Finally, Section 8 summarizes the main contributions of the paper and discusses current and future work.

## 2 CP-nets

CP-nets [3, 2] (for Conditional Preference nets) are a graphical model for compactly representing conditional and qualitative preference relations. They exploit conditional preferential independence by decomposing an agent's preferences via the **ceteris paribus** (cp) assumption. Informally, CP-nets are sets of ***ceteris paribus (cp)*** preference statements. For instance, the statement *"I prefer red wine to white wine if meat is served."* asserts that, given two meals that differ *only* in the kind of wine served *and* both containing meat, the meal with a red wine is preferable to the meal with a white wine. On the other hand, this statement does not order two meals with a different main course. Many users' preferences appear to be of this type.

CP-nets bear some similarity to Bayesian networks. Both utilize directed graphs where each node stands for a domain variable, and assume a set of *features* (variables) $F = \{X_1, \ldots, X_n\}$ with the corresponding finite domains $\mathcal{D}(X_1), \ldots, \mathcal{D}(X_n)$. For each feature $X_i$, a user specifies a (possibly empty) set of ***parent*** features $Pa(X_i)$ that can affect her preferences over the values of $X_i$. This defines a ***dependency graph*** in which each node $X_i$ has $Pa(X_i)$ as its immediate predecessors.

Given this structural information, the user explicitly specifies her preference over the values of $X_i$ for *each complete assignment* on $Pa(X_i)$. This preference is assumed to take the form of a linear ordering over $\mathcal{D}(X_i)$ [3, 2]. Each such specification is called below a ***preference statement*** for the variable $X_i$. These conditional preferences over the values of $X_i$ are captured by a *conditional preference table* which is annotated with the node $X_i$ in the CP-net. An ***outcome*** is an assignment of values to the variables with each value taken from the corresponding domain.

As an example, consider a CP-net whose features are $A$, $B$, $C$ and $D$, with binary domains containing $f$ and $\overline{f}$ if $F$ is the name of the feature, and with the following preference statements:

$d : a \succ \overline{a},\ \ \overline{d} : a \succ \overline{a},$

$a : b \succ \overline{b},\ \ \overline{a} : \overline{b} \succ b,$

$b : c \succ \overline{c},\ \ \overline{b} : \overline{c} \succ c,$

$c : d \succ \overline{d},\ \ \overline{c} : \overline{d} \succ d.$

Here the preference statement $d : a \succ \overline{a}$ states that $A = a$ is preferred to $A = \overline{a}$, given that $D = d$. From the structure of these preference statements we see that $Pa(A) = \{D\}, Pa(B) = \{A\}, Pa(C) = \{B\}, Pa(D) = \{C\}$ so the dependency graph is cyclic.

An ***acyclic*** CP-net is one in which the dependency graph is acyclic. As an example, consider a CP-net whose features and domains are as above and with the following preference statements:

$a \succ \overline{a},$

$b \succ \overline{b},$

$(a \wedge b) \vee (\overline{a} \wedge \overline{b}) : c \succ \overline{c},\ \ (a \wedge \overline{b}) \vee (\overline{a} \wedge b) : \overline{c} \succ c,$

$c : d \succ \overline{d},\ \ \overline{c} : \overline{d} \succ d.$

Here, the preference statement $a \succ \overline{a}$ represents the unconditional preference for $A = a$ over $A = \overline{a}$. Also each preference statement for the variable $C$ is a actually an abbreviated version of two preference statements. In this example we have $Pa(A) = \emptyset, Pa(B) = \emptyset, Pa(C) = \{A, B\}, Pa(D) = \{C\}$.

The semantics of CP-nets depends on the notion of a ***worsening flip***. A worsening flip is a transition between two outcomes that consists of a change in the value of a single variable to one which is less preferred in the unique preference statement for that variable. By analogy we define an ***improving flip***. For example, in the acyclic CP-net above, passing from $abcd$ to $ab\overline{c}d$ is a worsening flip since $c$ is better than $\overline{c}$ given $a$ and $b$. We say that an outcome $\alpha$ is ***better*** than the outcome $\beta$ (or, equivalently, $\beta$ is ***worse*** than $\alpha$), written as $\alpha \succ \beta$, iff there is a chain of worsening flips from $\alpha$ to $\beta$. This definition induces a strict preorder over the outcomes. In the above acyclic CP-net the outcome $\overline{a}b\overline{c}\overline{d}$ is worse than $abcd$.

An ***optimal*** outcome is one for which no better outcome exists. In general, a CP-net does not need to have an optimal outcome. As an example consider two features $A$ and $B$ with the respective domains $\{a, \overline{a}\}$ and $\{b, \overline{b}\}$ and the following preference statements:

$a : b \succ \overline{b},\ \ \overline{a} : \overline{b} \succ b,$

$b : \overline{a} \succ a,\ \ \overline{b} : a \succ \overline{a}.$

It is easy to see that then

$ab \succ a\overline{b} \succ \overline{a}\overline{b} \succ \overline{a}b \succ ab.$

## 3 Strategic games with parametrized preferences

In this section we introduce a generalization of the notion of a strategic game used in game theory, see, e.g., [5].

First we need the concept of a ***preference*** on a set $A$ which in this paper denotes a strict linear ordering on $A$. If $\succ$ is a preference, we denote by $\succeq$ the corresponding ***weak preference*** defined by: $a \succeq b$ iff $a \succ b$ or $a = b$.

Given a sequence of non-empty sets $S_1, \ldots, S_n$ and $s \in S_1 \times \ldots \times S_n$ we denote the $i$th element of $s$ by $s_i$ and use the following standard notation of game theory, where $I := i_1, \ldots, i_k$ is a subsequence of $1, \ldots, n$:

- $s_{-i} := (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$,
- $s_I := (s_{i_1}, \ldots, s_{i_k})$,
- $(s'_i, s_{-i}) := (s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_n)$, where we assume that $s'_i \in S_i$,
- $S_{-i} := S_1 \times \ldots \times S_{i-1} \times S_{i+1} \times \ldots \times S_n$,
- $S_I := S_{i_1} \times \ldots \times S_{i_k}$.

In game theory it is customary to study strategic games in which the outcomes are numerical values provided by means of the payoff functions. A notable exception is [6] in which instead of payoff functions the linear quasi-orderings on the sets of joint strategies are used.

In our setup we adopt a different approach according to which each player has to his disposal a strict preference relation $\succ(s_{-i})$ on his set of strategies *parametrized* by a joint strategy $s_{-i}$ of his opponents. So in our approach

- for each $i \in [1..n]$ player $i$ has a finite, non-empty, set $S_i$ of strategies available to him,
- for each $i \in [1..n]$ and $s_{-i} \in S_{-i}$ player $i$ has a preference relation $\succ(s_{-i})$ on his set of strategies $S_i$.

In what follows such a ***strategic game with parametrized preferences*** (in short a ***game with parametrized preferences***, or just a ***game***) for $n$ players is represented by a sequence

$$(S_1, \ldots, S_n, \succ(s_{-1}), \ldots, \succ(s_{-n})),$$

where each $s_{-i}$ ranges over $S_{-i}$.

It is straightforward to transfer to the case of games with parametrized preferences the basic notions concerning strategic games. In particular, given a game $G$ with parametrized preferences specified as above we say that a joint strategy $s$ is a (pure) ***Nash equilibrium*** of $G$ if for all $i \in [1..n]$ and all $s'_i \in S_i$

$$s_i \succeq(s_{-i})\ s'_i.$$

(For the original definition see, e.g., [5].)

To clarify this definitions consider the classical Prisoner's dilemma strategic game represented by the following bimatrix representing the payoffs to both players:

| | $C_2$ | $N_2$ |
|---|---|---|
| $C_1$ | $3, 3$ | $0, 4$ |
| $N_1$ | $4, 0$ | $1, 1$ |

So each player $i$ has two strategies, $C_i$ (cooperate) and $N_i$ (not cooperate), the payoff to player 1 for the joint strategy $(C_1, N_2)$ is 0, etc. To represent this game as a game with parametrized preferences we simply stipulate that

$\succ(C_2) := N_1 \succ C_1$, $\succ(N_2) := N_1 \succ C_1$,
$\succ(C_1) := N_2 \succ C_2$, $\succ(N_1) := N_2 \succ C_2$.

These orderings reflect the fact that for each strategy of the opponent each player considers his 'not cooperate' strategy better than his 'cooperate' strategy. It is easy to check that $(N_1, N_2)$ is a unique Nash equilibrium of this game with parametrized preferences.

## 4 From CP-nets to strategic games

Consider now a CP-net with the set of variables $\{X_1, \ldots, X_n\}$ with the corresponding finite domains $\mathcal{D}(X_1), \ldots, \mathcal{D}(X_n)$. We write each preference statement for the variable $X_i$ as $X_I = a_I : \succ_i$, where for the subsequence $I = i_1, \ldots, i_k$ of $1, \ldots, n$:

- $Pa(X_i) = \{X_{i_1}, \ldots, X_{i_k}\}$,
- $X_I = a_I$ is an abbreviation for $X_{i_1} = a_{i_1} \wedge \ldots \wedge X_{i_k} = a_{i_k}$,
- $\succ_i$ is a preference over $\mathcal{D}(X_i)$.

We also abbreviate $\mathcal{D}(X_{i_1}) \times \ldots \times \mathcal{D}(X_{i_k})$ to $\mathcal{D}(X_I)$.

By definition, the preference statements for a variable $X_i$ are exactly all statements of the form $X_I = a_I : \succ(a_I)$, where $a_I$ ranges over $\mathcal{D}(X_I)$ and $\succ(a_I)$ is a preference on $\mathcal{D}(X_i)$ that depends on $a_I$.

We now associate with each CP-net $N$ a game $\mathcal{G}(N)$ with parametrized preferences as follows:

- each variable $X_i$ corresponds to a player $i$,
- the strategies of player $i$ are the elements of the domain $\mathcal{D}(X_i)$ of $X_i$.

To define the parametrized preferences, consider a player $i$. Suppose $Pa(X_i) = \{X_{i_1}, \ldots, X_{i_k}\}$ and let $I := i_1, \ldots, i_k$. So $I$ is a subsequence of $1, \ldots, i-1, i+1, \ldots, n$. Given a joint strategy $a_{-i}$ of the opponents of player $i$, we associate with it the preference relation $\succ(a_I)$ on $\mathcal{D}(X_i)$ where $X_I = a_I : \succ(a_I)$ is the unique preference statement for $X_i$ determined by $a_I$.

In words, the preference of a player $i$ over his strategies, given the strategies chosen by its opponents, say $a_{-i}$, coincides with the preference given by the CP-net over the domain of $X_i$ given the assignment to his parents $a_I$ which must coincide with the projection of $a_{-i}$ over $I$. This completes the definition of $\mathcal{G}(N)$.

As an example consider the first CP-net of Section 2. The corresponding game has four players $A$, $B$, $C$, $D$, each with two strategies indicated with $f$, $\bar{f}$ for player $F$. The preference of each player on his strategies will depend only on the strategies chosen by the players which correspond to his parents in the CP-net. Consider for example player $B$. His preference over his strategies $b$ and $\bar{b}$, given the joint strategy of his opponents $s_{-B} = dac$, is $b \succ \bar{b}$. Notice that, for example, the same ordering holds for the opponents joint strategy $s_{-B} = \bar{d}a\bar{c}$, since the strategy chosen by the only player corresponding to his parent, $A$, has not changed.

We have then the following result.

**Theorem 1** *An outcome of a CP-net $N$ is optimal iff it is a Nash equilibrium of the game $\mathcal{G}(N)$.*

## 5 From strategic games to CP-nets

We now associate with each game $G$ with parametrized preferences a CP-net $\mathcal{N}(G)$ as follows:

- each player $i$ corresponds to a variable $X_i$,
- the domain $\mathcal{D}(X_i)$ of the variable $X_i$ consists of the set of strategies of player $i$,
- we stipulate that $Pa(X_i) = \{X_1, X_{i-1}, \ldots, X_{i+1}, \ldots, X_n\}$, where $n$ is the number of players in $G$.

Next, for each joint strategy $s_{-i}$ of the opponents of player $i$ we take the preference statement $X_{-i} = s_{-i} : \succ(s_{-i})$, where $\succ(s_{-i})$ is the preference relation on the set of strategies of player $i$ associated with $s_{-i}$.

This completes the definition of $\mathcal{N}(G)$. As an example of this construction let us return to the Prisoner's dilemma game with parametrized preferences from Section 3. In the corresponding CP-net we have then two variables $X_1$ and $X_2$ corresponding to players 1 and 2, with the respective domains $\{C_1, N_1\}$ and $\{C_2, N_2\}$. To explain how each parametrized preference translates to a preference statement take for example $\succ(C_2) := \ N_1 \succ C_1$. It translates to $X_2 = C_2 : \ N_1 \succ C_1$.
We have now the following counterpart of Theorem 1.

**Theorem 2** *A joint strategy is a Nash equilibrium of the game $G$ iff it is an optimal outcome of the CP-net $\mathcal{N}(G)$.*

## 6 Reduced CP-nets

The disadvantage of the above construction of the CP-net $\mathcal{N}(G)$ from a game $G$ is that it always produces a CP-net in which all sets of parent features are of size $n-1$ where $n$ is the number of features of the CP-net. This can be rectified by reducing each set of parent features to a minimal one as follows.

Given a CP-net $N$, consider a variable $X_i$ with the parents $Pa(X_i)$, and take a variable $Y \in Pa(X_i)$. Suppose that for all assignments $a$ to $Pa(X) - \{Y\}$ and any two values $y_1, y_2 \in \mathcal{D}(Y)$, the orderings $\succ(a, y_1)$ and $\succ(a, y_2)$ on $\mathcal{D}(X_i)$ coincide.

We say then that $Y$ is *redundant* in the set of parents of $X_i$. It is easy to see that by removing all redundant variables from the set of parents of $X_i$ and by modifying the corresponding preference statements for $X_i$ accordingly, the strict preorder $\succ$ over the outcomes of the CP-nets is not changed.

Given a CP-net, if for all its variable $X_i$ the set $Pa(X_i)$ does not contain any redundant variable, we say that the CP-net is ***reduced***.

By iterating the above construction every CP-net can be transformed to a reduced CP-net. As an example consider a CP-net with three features, $X, Y$ and $Z$, with the respective domains $\{a_1, a_2\}, \{b_1, b_2\}$ and $\{c_1, c_2\}$. Suppose now that $Pa(X) = Pa(Y) = \emptyset, Pa(Z) = \{X, Y\}$ and that
$\succ(a_1, b_1) = \succ(a_2, b_1), \ \succ(a_1, b_2) = \succ(a_2, b_2),$
$\succ(a_1, b_1) = \succ(a_1, b_2), \ \succ(a_2, b_1) = \succ(a_2, b_2).$
Then both $X$ and $Y$ are redundant, so we can reduce the CP-net by reducing $Pa(Z)$ to $\emptyset$. $Z$ becomes an independent variable in the reduced CP-net with an ordering over its domain which coincides with the unique one given in the original CP-net in terms of the assignments to its parents.

In what follows for a CP-net $N$ we denote by $r(N)$ the corresponding reduced CP-net. The following result summarizes the relevant properties of $r(N)$ and relates it to the constructions of $\mathcal{G}(N)$ and $\mathcal{N}(G)$.

**Theorem 3**

*(i) Each CP-net $N$ and its reduced form $N' = r(N)$ have the same ordering $\succ$ over the outcomes.*

*(ii) For each CP-net $N$ and its reduced form $N' = r(N)$ we have $\mathcal{G}(N) = \mathcal{G}(N')$.*

*(iii) Each reduced CP-net $N$ is a reduced CP-net corresponding to the game $\mathcal{G}(N)$. Formally: $N = r(\mathcal{N}(\mathcal{G}(N)))$.*

Part $(i)$ states that the reduction procedure preserves the ordering over the outcomes. Part $(ii)$ states that the construction of a game corresponding to a CP-net does not depend on the redundancy of the given CP-net. Finally, part $(iii)$ states that the reduced CP-net $N$ can be obtained 'back' from the game $\mathcal{G}(N)$.

Games $G$ such that the CP-net $N' = r(\mathcal{N}(G))$ is acyclic are not uncommon. In fact, they naturally represent multi-agent scenarios where agents (that is, players of the game) can be partitioned into levels $1, 2, \ldots, n$, such that agents at level $i$ can express their preferences (that is, payoff function) without looking at what players at higher levels do. Informally, agents at level $i$ are more important than agents at level $j$ is $j > i$. In particular, agents at level 1 can decide their preferences without looking at the behavior of any other agent.

## 7 Game-theoretic techniques in CP-nets

Given the correspondence between CP-nets and games and its properties presented in the previous sections, we can now use them to transfer standard techniques of game theory, used to find Nash equilibria, to CP-nets to find their optimal outcomes.

More specifically, we can transfer two techniques of iterated elimination of 'suboptimal' strategies —those that are strictly dominated or are never best responses (see, e.g., [6].) To introduce them in the context of CP-nets consider a CP-net $N$ with the set of variables $\{X_1, \ldots, X_n\}$ with the corresponding finite domains $\mathcal{D}(X_1), \ldots, \mathcal{D}(X_n)$.

- We say that an element $d_i$ from the domain $\mathcal{D}(X_i)$ of the variable $X_i$ is a ***best response*** to a preference statement

$$X_I = a_I : \succ_i$$

for $X_i$ if $d_i \succeq_i d'_i$ for all $d'_i \in \mathcal{D}(X_i)$.

- We say that an element $d_i$ from the domain of the variable $X_i$ is a ***never a best response*** if it is not a best response to any preference statement for $X_i$.

- Given two elements $d_i, d'_i$ from the domain $\mathcal{D}(X_i)$ of the variable $X_i$ we say that $d'_i$ is ***strictly dominated*** by $d_i$ if for all preference statements $X_I = a_I : \succ_i$ for $X_i$ we have

$$d_i \succ_i d'_i.$$

By a ***subnet*** of a CP-net $N$ we mean a CP-net obtained from $N$ by removing some elements from some variable domains followed by the removal of all preference statements that refer to a removed element.

Then we introduce the following relation between a CP-net $N$ and its subnet $N'$:

$$N \rightarrow_{NBR} N'$$

when $N \neq N'$ and for each variable $X_i$ each removed element from the domain of $X_i$ is never a best response in $N$, and introduce an analogous relation $N \rightarrow_S N'$ for the case of strictly dominated elements.

The following result then holds.

**Theorem 4** *Suppose that $N \rightarrow^*_{NBR} N'$, i.e., the CP-net $N'$ is obtained by an iterated elimination of never best responses from the CP-net $N$.*

*(i) Then $s$ is an optimal outcome of $N$ iff it is an optimal outcome of $N'$.*

*(ii) If each variable in $N'$ has a singleton domain, then the resulting outcome is a unique optimal outcome of $N$.*

*(iii) All iterated eliminations of never best responses from the CP-net $N$ yield the same final outcome.*

To illustrate the use of this theorem reconsider the first CP-net from Section 2, i.e., the one with the preference statements

$d : a \succ \overline{a}$, $\overline{d} : a \succ \overline{a}$,
$a : b \succ \overline{b}$, $\overline{a} : \overline{b} \succ b$,
$b : c \succ \overline{c}$, $\overline{b} : \overline{c} \succ c$,
$c : d \succ \overline{d}$, $\overline{c} : \overline{d} \succ d$.

Denote it by $N$.

We can reason about it using the iterated elimination of strictly dominated strategies (which coincides here with the iterated elimination of never best responses, since each domain has exactly two elements).

We have the following chain of reductions:

$$N \rightarrow_S N_1 \rightarrow_S N_2 \rightarrow_S N_3 \rightarrow_S N_4,$$

where

- $N_1$ results from $N$ by removing $\overline{a}$ (from the domain of $A$) and the preference statements $d : a \succ \overline{a}$, $\overline{d} : a \succ \overline{a}$, $\overline{a} : \overline{b} \succ b$,
- $N_2$ results from $N_1$ by removing $\overline{b}$ and the preference statements $a : b \succ \overline{b}$, $\overline{b} : \overline{c} \succ c$,
- $N_3$ results from $N_2$ by removing $\overline{c}$ and the preference statements $b : c \succ \overline{c}$ $\overline{c} : \overline{d} \succ d$,
- $N_4$ results from $N_3$ by removing $\overline{d}$ from the domain of $D$ and the preference statement $c : d \succ \overline{d}$.

Indeed, in each step the removed element is strictly dominated in the considered CP-net. So using the iterated elimination of strictly dominated elements we reduced the original CP-net to one in which each variable has a singleton domain and consequently found a unique optimal outcome of the original CP-net $N$.

Finally, the following result shows that the introduced reduction relation on CP-nets is complete for acyclic CP-nets.

**Theorem 5** *For each acyclic CP-net $N$ a unique subnet $N'$ with the singleton domains exists such that $N \rightarrow^*_{NBR} N'$.*

## 8 Conclusions and future work

In this paper we related two formalisms that are commonly used in reasoning about multi-agent systems, strategic games and CP-nets. To this end we generalized the concept of strategic games to games with parametrized preferences and showed that optimal outcomes in CP-nets are exactly Nash equilibria of such games. This allowed us to exploit game-theoretic techniques in search for optimal outcomes of CP-nets.

Our current research deals two other aspects concerning strategic games and preferences. First, thanks to the established correspondence, we can also use the techniques developed to reason about optimal outcomes of a CP-net in search for Nash equilibria of strategic games with parametrized preferences. These techniques, as recently shown in [4, 7], involve the use of the customary constraint solving techniques. In fact, it has been shown that the optimal outcomes of any CP-net, even a cyclic one, can be found by just solving a set of hard constraints. Thus hard constraint solving is enough to find also Nash equilibria in strategic games.

Second, we found that the direct correpondence between the optimal solutions of a CP-net and the Nash equilibria of the corresponding game cannot be easily found in other preference modelling formalisms, for example soft constraints, see [1]. In fact, it is possible to show that, in a so-called fuzzy constraint problem, there can be optimal solutions which are not Nash equilibria of corresponding games, and vice-versa. We are therefore studying the conditions under which soft constraints can be related to game theory.

In this paper we assumed that payoff functions give a linear order over the strategies of a player. It could be useful to see whether our results can be generalized to games in which players' strategies can be incomparable or indifferent to each other, thus using partial orderings with ties. We are currently studying this scenario.

This paper is just a first step towards what we think is a fruitful cross-fertilization between preferences, constraint solving, and game theory.

## References


[1] S. Bistarelli, U. Montanari, and F. Rossi. Semiring-based constraint solving and optimization. *Journal of the ACM*, 44(2):201–236, Mar 1997.

[2] C. Boutilier, R. I. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.

[3] C. Boutilier, R. I. Brafman, H. H. Hoos, and D. Poole. Reasoning with conditional ceteris paribus preference statements. In *UAI '99*, pages 71–80. Morgan Kaufmann, 1999.

[4] R. Brafman and Y. Dimopoulos. Extended semantics and optimization algorithms for cp-networks. *Computational Intelligence*, 20, 2:218–245, 2004.

[5] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard Univ Press, Cambridge, Massachusetts, 1991.

[6] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.

[7] S. Prestwich, F. Rossi, K. B. Venable, and T. Walsh. Constraint-based preferential optimization. In *Proc. of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*. Morgan Kaufmann, 2005.